Describing Human Aesthetic Perception by Deeply-learned Attributes from Flickr

Luming Zhang

arXiv:1605.07699v1 [cs.CV] 25 May 2016

***Abstract*—Many aesthetic models in computer vision suffer from two shortcomings: 1) the low descriptiveness and interpretability of those hand-crafted aesthetic criteria (*i.e.*, non-indicative of region-level aesthetics), and 2) the difficulty of engineering aesthetic features adaptively and automatically toward different image sets. To remedy these problems, we develop a deep architecture to learn aesthetically-relevant visual attributes from Flickr[1], which are localized by multiple textual attributes in a weakly-supervised setting. More specifically, using a bag-of-words (BoW) representation of the frequent Flickr image tags, a sparsity-constrained subspace algorithm discovers a compact set of textual attributes (*e.g.*, landscape and sunset) for each image. Then, a weakly-supervised learning algorithm projects the textual attributes at image-level to the highly-responsive image patches at pixel-level. These patches indicate where humans look at appealing regions with respect to each textual attribute, which are employed to learn the visual attributes. Psychological and anatomical studies have shown that humans perceive visual concepts hierarchically. Hence, we normalize these patches and feed them into a five-layer convolutional neural network (CNN) to mimick the hierarchy of human perceiving the visual attributes. We apply the learned deep features on image retargeting, aesthetics ranking, and retrieval. Both subjective and objective experimental results thoroughly demonstrate the competitiveness of our approach.**

***Index Terms*—Deep architecture, Aesthetics, Flickr, Attributes, Convolutional neural network**

I. Introduction

Perceptually aesthetic modeling refers to the process of discovering low-level and high-level visual patterns that can arouse human aesthetic perception. It is a useful technique to enhance many applications, such as image retargeting, photo album management and scene rendering. Take retargeting for example, the aesthetically pleasing image regions are squeezed slightly and vice versa. Moreover, effectively describing region-level aesthetics can guide the level of details in non-photorealistic image rendering. As a popular media sharing website with millions of photos uploaded and commented daily, Flickr can be considered as an ideal platform to study and simulate how humans perceive photos with aesthetic features at both low-level and high-level. However, perceptually aesthetic modeling based on Flickr will encounter the following challenges:

- Flickr contains a large amount of photos with multiple latent aesthetic attributes (*e.g.*, the "field guide" and "movement"). Conventional aesthetic models, however, are typically built upon generic features which quantify the compliance to a pre-defined criterion (*e.g.*, the "rule of thirds"). In practice, we need tailored and dataset-dependent features to capture these latent aesthetic attributes, but engineering these features may require the domain knowledge of professional photographers.
- The interpretability of an aesthetic model means the capability of indicating which regions are responsive to each aesthetic attribute. A well interpretable aesthetic model has both scientific impact and application value. As far as we know, the existing aesthetic modeling pipeline is more or less a blackbox. They have limited power to provide region-level response to each aesthetic attribute, particularly the abstract and global ones such as "harmonic", "dream", and "colorful".
- Previous methods associate each photo with a set of aesthetic attributes in well-controlled settings. For example, the training images' beautifulness is quantified to reduce the disturbance of noisy tags [2]. In reality, image tags from Flickr are labeled uncontrollably by a variety of users with diverse backgrounds, education, *etc*. This will inevitably produces noisy image tags, which necessitate an aesthetic model to robustly handle noises from image tags.

To solve these problems, we propose a CNN-based framework which models aesthetic perception automatically and interpretably, by calculating a compact set of textual and visual attributes from tagged Flickr images. The pipeline of our proposed framework is presented in Fig. 1. For each tagged image from a large-scale Flickr corpus, we employ a bag-of-words (BoW) representation of those frequently-occurring and noisy image tags. Toward an efficient system, a sparsity-constrained subspace algorithm projects each BoW histogram onto a compact set of textual attributes. They reflect the highly representative aesthetic attributes of an image. To locate the appealing visual attributes, a weakly supervised algorithm maps the textual attributes at image-level onto the salient patches (indicated by different colors in the middle of Fig. 1) in an image. Both psychological and anatomical studies have shown that human vision system is multi-layered and forms higher-level abstracts from input raw pixels incrementally. That means the intrinsic hierarchical structure of CNN has the potential to model human visual perception. Based on this, for each textual attribute, the corresponding extracted patches are employed to learn a five-layer CNN to simulate the hierarchial perception of human beings. The learned deep feature is applied on a range of multimedia tasks: image retargeting, aesthetics classification and image retrieval.

L. Zhang is with the School of Computing, National University Singapore.

[1]https://www.flickr.com/

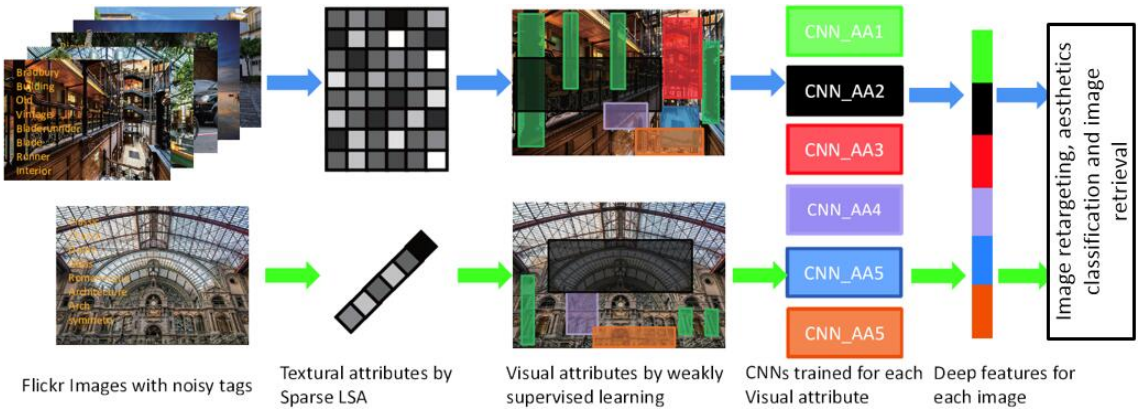

Fig. 1. The pipeline of the proposed CNN-based aesthetic modeling framework (The blue and green arrows denote the training and test phases respectively. The color-tagged regions indicate visual attributes which are aesthetically pleasing.).

The contributions of this paper can be summarized as: 1) The first deep architecture that learns aesthetically-relevant visual attributes extracted from massive-scale Flickr images; 2) a weakly supervised algorithm associating each textual attribute with the corresponding visual attribute; and 3) adopting the deep features on several multimedia applications, combined with extensive experimental validation.

The rest of this paper is organized as follows: Sec II briefly reviews the related work. Sec III introduces the proposed deep image aesthetics modeling framework, including the sparsity-constrained textual attributes mining, weakly supervised visual attributes detection, our developed CNN and it applications. Experimental results in Sec IV thoroughly demonstrate the effectiveness of our method. Sec V concludes.

II. Related Work

Our work is related to two research topics in multimedia domain [50], [51]: computational image aesthetics analysis and deep learning-based aesthetic attribute modeling.

A. Computational Image Aesthetic Analysis

1) Global features-based models: Datta *et al.* [3] proposed 58 low-level visual features, *e.g.*, the shape convexity, to capture photo aesthetics. Dhar *et al.* [5] proposed a set of high-level attribute-based predictors to evaluate photo aesthetics. In [6], Luo *et al.* employed a GMM-based hue distribution and a prominent line-based texture distribution to assess the global composition of an image. To represent image local composition, regional features describing human faces, region clarity and region complexity were developed. In [7], Marchesotti *et al.* tackled the problem of visual aesthetic modeling by discovering mid-level features. The designed algorithm can automatically learn semantic aesthetics-related attributes by combining image, scoring, and textual data from the AVA data set. The authors shown that the learned attributes can facilitate a variety of media applications, *e.g.*, aesthetic quality prediction, image tagging and retrieval. Experimental results shown that the above two generic descriptors outperform a variety of hand-crafted and dataset-dependent aesthetic descriptors.

2) Local feature-based models: Cheng *et al.* [36] proposed the omni-range context, *i.e.*, the spatial distribution of arbitrary pairwise image patches, to describe image composition. Nishiyama [9] *et al.* first detected multiple subject regions in a photo, and afterward an SVM classifier is trained for each subject region. Finally, the aesthetics of an image is quantified by combining the SVM scores of a photo's internal subject regions. In [10], Nishiyama *et al.* proposed a color harmony-based aesthetic model, which models image color distribution by leveraging its patches. The patch-level color distribution is integrated into a BoW histogram, which is classified by an SVM to determine whether a photo is highly or low aesthetic. Bhattacharya *et al.* [11] developed a spatial recomposition system which allows users to select a foreground object interactively. The system presents recommendations to indicate an optimal location of the foreground object, which is detected by integrating multiple visual features.

B. Deep learning-based Aesthetic Attribute Modeling

As far as we know, there are only two deep learning models for visual aesthetic analysis. In [22], Lu *et al.* proposed a double-layer CNN architecture to automatically discover effective features that capture image aesthetics from two heterogeneous input sources, i.e., aesthetic features from both the global and local views. The double CNN layers are jointly trained from two inputs. The first layer takes global image representation as the input, while the second layer takes local image representations as the input. This allows us to leverage both compositional and local visual information. Based on the evaluation from the AVA data set, Lu *et al.*'s algorithms significantly outperforms the results reported earlier. This model differs from ours fundamentally in two aspects: 1) Both the global and local views are heuristically defined, there is no guarantee that they can well locate the aesthetically pleasing regions across different datasets. Comparatively, our approach uses a weakly supervised algorithm to discover visually appealing regions indicated by tags. Thus, it can be conveniently adapted onto different datasets; 2) Lu *et al.*'s model simply captures the global and local aesthetic features of a photo. But there is no evidence that abstract aesthetic cues such as "vivid" and "harmonic" can be well described. Noticeably, in our

model, a set of CNNs are trained. Each encodes visual attribute corresponding to a textual attribute, which can capture either a concrete or abstract aesthetic cue. In [23], Champbell *et al.* trained two Restricted Boltzmann Machines (RBMs) on highly and low aesthetic images respectively. The authors observed that 10% of the filters learned from the highly aesthetic images capture the aesthetics-relevant visual cues. But this model is only available on simple abstract paintings with low resolution. It is computationally intractable to describe high resolution Flickr images with various semantics.

III. The Proposed Method

A. Sparse Textual Attributes Discovery

Given a Flickr image, we use an M-dimensional augmented frequency vector $\overrightarrow{\alpha}$ to represent the distribution of its tags[2]. In particular, to avoid the randomly-occurring noisy image tags, we select the M most frequent tags from the training image set. Then, we treat the tag set of each Flickr image as a document $\mathcal{D}$, based on which the i-th element of vector $\overrightarrow{\alpha}$ can be calculated as:

$$\overrightarrow{\alpha}(i) = 0.5 + \frac{0.5 * f(i, \mathcal{D})}{\max(f(j, \mathcal{D}) : j \in \mathcal{D})}, \tag{1}$$

where $f(i, \mathcal{D})$ counts the times of the i-th tag from the M most frequent ones occurring in document d, and the denominator functions as a normalization factor. In our implementation, we set $M = 100$ based on cross validation.

Given N Flickr images, they can be represented as N augmented frequency vectors. Thereafter, we column-wise stack them into a matrix $\mathbf{X} = [\overrightarrow{\alpha}_1; \overrightarrow{\alpha}_2; \cdots; \overrightarrow{\alpha}_N] \in \mathbb{R}^{N \times M}$, where each row $\vec{X}_j \in \mathbb{R}^N$ denotes the j-th feature vector cross all the documents. To obtain the textual attributes of each Flickr image, we adopt a subspace algorithm, which converts the original M-dimensional vector corresponding to each Flickr image into a D-dimensional textual attribute vector ($D < \min(M, N)$). In our implementation, all the 75246 training images are adopted to create document matrix $\mathbf{X}$

Following the latent semantic analysis (LSA) [25], we assume that the D textual attributes $\{\vec{u}_1, \vec{u}_2, \cdots, \vec{u}_D\}$ are uncorrelated, where each attribute $\vec{u}_d \in \mathbb{R}^N$ has the unit length, *i.e.*, $\|\vec{u}_d\|_2 = 1$. Denote $\mathbf{U} = [\vec{u}_1, \vec{u}_2, \cdots, \vec{u}_D] \in \mathbb{R}^{N \times D}$, we have $\mathbf{U}^T\mathbf{U} = \mathbf{I}$ where $\mathbf{I}$ is the identity matrix. It is reasonable to assume that each feature vector $\vec{X}_j$ can be linearly reconstructed by the textual attributes:

$$\mathbf{X}_j = \sum\nolimits_{d=1}^{D} a_{dj}\mathbf{u}_d + \epsilon_j, \tag{2}$$

In the matrix form, the above equation can be reorganized into $\mathbf{X} = \mathbf{U}\mathbf{A} + \epsilon$, where $\mathbf{A} = [a_{dj}] \in \mathbb{R}^{D \times M}$ denotes the projection matrix from the tag space to the textual attribute space. We can obtain the projection matrix $\mathbf{A}$ by solving the following optimization. It minimizes the rank-D approximation error subject to the orthogonality constraint of $\mathbf{U}$:

$$\min_{\mathbf{U},\mathbf{A}} \|\mathbf{X} - \mathbf{U}\mathbf{A}\|_F^2, \quad s.t. \quad \mathbf{U}^T\mathbf{U} = \mathbf{I}, \tag{3}$$

where $\|\cdot\|_F$ denotes the matrix Frobenius norm. The constraint $\mathbf{U}\mathbf{U}^T = \mathbf{I}$ reflects the uncorrelated property among textual attributes.

Sparsity of textual attributes: The projection matrix $\mathbf{U}$ learned from (3) can reconstruct matrix $\mathbf{X}$ by a linear combination of textual attributes. Typically, the number of textual attributes D is not small[3]. It means that aesthetically modeling an Flickr image by analyzing all its correlated textual attributes might be intractable. Toward an efficient aesthetic model, we encode a sparsity constraint into (3) in order to achieve a sparse projection matrix $\mathbf{A}$. An entry-wise l_1-norm of $\mathbf{A}$ is added as a regularization term to the loss function. Based on this, we formulate the sparse textual attributes discovery as:

$$\min_{\mathbf{U},\mathbf{A}} \|\mathbf{X} - \mathbf{U}\mathbf{A}\|_F^2 + \lambda\|\mathbf{A}\|_1, \quad s.t., \quad \mathbf{U}^T\mathbf{U} = \mathbf{I}, \tag{4}$$

where $\|\mathbf{A}\|_1 = \sum_{d=1}^{D}\sum_{j=1}^{M}|a_{dj}|$ is the entry-wise l_1-norm of $\mathbf{A}$; and λ is the positive regularization parameter controlling the density of $\mathbf{A}$, *i.e.*, the number of nonzero entries. In general, a larger λ leads to a sparser $\mathbf{A}$. On the other hand, a too sparse $\mathbf{A}$ will lose some relationships between Flickr tags and textual attributes and will in turn harm the reconstruction quality. In practice, it is important to select an appropriate λ to obtain a more sparse $\mathbf{A}$ while still achieving a good reconstruction performance. It is noticeable that (4) is solved by alternatively optimizing matrix $\mathbf{U}$ and $\mathbf{A}$, as detailed in [26].

B. Weakly Supervised Visual Attributes Learning

This section introduces a graphlet-based weakly supervised learning framework to detect visual attributes corresponding to each textual attribute. We first construct a superpixel pyramid which captures objects with different shapes seamlessly. To quantify region-level response to each textual attribute, a manifold embedding algorithm transfers textual attributes into graphlets. Finally, the patch containing graphlet most responsive to each textural attribute form the visual attribute.

Graphlet construction: There are usually millions of raw pixels in an image, treating each of them independently brings intractable computation. It is generally accepted that pixels within an images are highly correlated with its spatial neighbors. Therefore, we sample a collection of superpixels and use them to construct different objects. As objects with different scales may evoke human aesthetic perception, we construct superpixels with three sizes by overly, moderately and deficiently segmenting each image. The segmentation is based on the well-known simple linear iterative clustering (SLIC). It is built upon the k-means clustering and has a time complexity of $O(N)$, where N is the number of pixels in an image. Compared with the conventional methods, experiments shown that SLIC is computationally more efficient, requires less memory and generates superpixels more adherent to object boundaries. By segmenting each image three times with different SLIC parameters, a superpixel pyramid is constructed to capture objects with different sizes.

As shown in Fig. 2, different objects can be constructed by a set of spatially neighboring superpixels. More specifically, a

[2]In our implementation, we set $M = 100$.

[3]As experimented on our own complied data set, each Flickr image is associated with $2 \sim 5$ textural attributes and the optimal total number of attributes D is between 10 and 15.

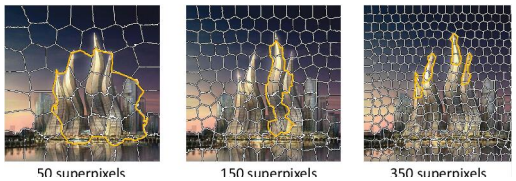

Fig. 2. A superpixel pyramid describing the Dubai towers from coarse to fine. The first pyramid layer describes the rough outline of the four towers, the surrounding lake and the sky, reflecting textual attributes such as "harmonic" and "symmetric". The second layer captures the shape of each individual tower, which corresponds to textual attributes such as "flame" and "dynamic". The third layer encodes the details of the four towers, *e.g.*, the shape tower top. It is highly responsive to textual attributes such as "shape" and "edge".

graphlet is a moderately sized graph defined as:

$$G = (V, E), \tag{5}$$

where V is a set of vertices, each representing a superpixel; E is a set of edges, each connecting a pair of spatially neighboring superpixels. We call a graphlet t-sized if it is constituted by t superpixels.

Given a t-sized graphlet, we represent it by a $t\times(t+128+9)$ matrix as:

$$\mathbf{M} = [\mathbf{M}_C, \mathbf{M}_T, \mathbf{M}_S], \tag{6}$$

where $\mathbf{M}_C$ is a $t\times 9$ matrix and each row is the 9-dimensional color moment [29] from a superpixel; $\mathbf{M}_T$ is a $t\times 128$ matrix where each row is a 128-dimensional HOG [30] from a superpixel; and $\mathbf{M}_S$ is a $t\times t$ adjacency matrix representing the topology of a graphlet. The graphlet extraction is based on random walk on the superpixel mosaic. More specifically, we first index the superpixels and then select a starting one. Afterward, spatially neighboring superpixels are visited one-by-one until the maximum graphlet size is reached. Due to the number of graphlets is exponentially increasing the the maximum graphlet size, we set it to 7 in our implementation.

Weakly supervised semantic encoding: The textual attributes indicate the aesthetically pleasing regions in a Flickr image. To locate them, we propose a weakly supervised learning algorithm which transfers textual attributes into different graphlets in an image. The objective function is:

$$\arg\min_{\mathbf{Y}} \sum_{ij} \|y_i - y_j\|^2 [l_s(i,j) - l_d(i,j)], \;\; s.t., \;\; \mathbf{Y}\mathbf{Y}^T = \mathbf{I}_d, \tag{7}$$

where $\mathbf{Y} = [y_1, y_2, \cdots, y_n]$, each denoting a d-dimensional post-embedding graphlet from the training images. $\mathbf{Y}$ in Eq.(7) is initialized by the the vector obtaining by row-wise concatenating matix $\mathbf{M}$ in Eq.(6). The number of graphlets used in $\mathbf{Y}$ depends on different image sets. In our experiment, the number is 75246. l_s and l_d are functions measuring the semantic similarity and difference between graphlets, which are quantified according to textual attributes. Denoting $\overrightarrow{n} = [n^1, n^2, \cdots, n^C]^T$ where n^i is the number of images with the i-th textual attribute, and $c(\cdot)$ contains the textual attributes of the Flickr image from which a graphlet is extracted, then we can obtain:

$$l_s(i,j) = \frac{[c(G_i) \cap c(G_j)]\overrightarrow{n}}{\sum_c n^c}, \tag{8}$$

$$l_d(i,j) = \frac{[c(G_i) \oplus c(G_j)]\overrightarrow{n}}{\sum_c n^c}, \tag{9}$$

where the numerator of l_s denotes the number of images sharing the common textual attributes with the images where the i-th and j-th graphlets are extracted; the numerator of l_d is the number of images having different textual attributes with the images where the i-th and j-th graphlets are extracted. The denominator represents the total number of images with all textual attributes, which is used as a normalization factor. An example of calculating l_s and l_d is presented in Fig. 3, where $l_s(i,j) = \frac{5+3+7}{5+1+\cdots+1} = 0.333$ and $l_d(i,j) = \frac{1+3+1+4+3}{5+1+\cdots+1} = 0.200$.

Objective function (7) can be reorganized into:

$$\begin{aligned} &\arg\min_{\mathbf{Y}} \sum_{ij} \|y_i - y_j\|^2 [l_s(i,j) - l_d(i,j)] \\ &= \arg\max_{\mathbf{Y}} \mathrm{tr}(\mathbf{Y}\mathbf{U}\mathbf{Y}^T), \quad s.t. \;\; \mathbf{Y}\mathbf{Y}^T = \mathbf{I} \end{aligned} \tag{10}$$

where $\mathbf{U} = [-\overrightarrow{e}_{N-1}^T, \mathbf{I}_{N-1}]^T \mathbf{W}_1 [-\overrightarrow{e}_{N-1}^T, \mathbf{I}_{N-1}] + \cdots + [\mathbf{I}_{N-1}, -\overrightarrow{e}_{N-1}^T]^T \mathbf{W}_N [\mathbf{I}_{N-1}, -\overrightarrow{e}_{N-1}^T]$, and $\mathbf{W}_1$ is an $N\times N$ diagonal matrix whose h-th diagonal element is $l_s(h,i) - l_d(h,i)$. Note that (10) is a quadratic problem with quadratic constraints, which can be solved by eigenvalue decomposition with a time complexity of $O(N^3)$. To speedup the solution, an iterative coordinate propagation [31]-based embedding is applied.

Based on the manifold embedding algorithm, we calculate

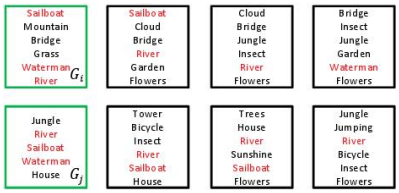

Fig. 3. An example illustrating the calculation of l_s and l_d from eight tagged Flickr images. The two green frames denote images where graphlet G_i and G_j are extracted. In this example, $c(G_i) \cap c(G_j)$ = {Sailboat, Waterman, River}, $c(G_i) \oplus c(G_j)$ = {Mountain, Bridge, Grass, Jungle, House}, and ={Sailboat: 5, Mountain: 1, Bridge: 3, Grass: 1, Waterman: 3, River: 7, Cloud: 2, Garden: 2, Flowers: 5, Jungle: 4, Insect: 4, House: 3, Tower: 1, Bicycle: 2, Sunshine: 1, Jumping: 1}.

the attribute-level response map for each image. In particular, given C the number of textual attributes, we train a multi-class SVM based on the one-versus-rest scenario. The saliency response of graphlet G to textual attribute $\mathbf{u}_d$ is calculated based on the probabilistic output of SVM:

$$p(y(G) \to \mathbf{u}_d) = \frac{1}{1 + \exp(-\phi_{\mathbf{u}_d}(y(G)))}, \tag{11}$$

where $y(G)$ is the post-embedding vector corresponding to graphlet G, and $\phi_{\mathbf{u}_d}(\cdot)$ denotes the classification hyper-plane that separates graphlets belonging to images with textual attribute $\mathbf{u}_d$ from the remaining graphlets.

After obtaining the saliency map to each textual attribute, we randomly generate 10000 patches in each image. The patch sizes are tuned as follows: the patch width is tuned from $[0.1w, 0.9w]$ with a step of $0.01w$, while the patch height is tuned from $[0.1h, 0.9h]$ with a step of $0.01h$. Note that w and h denote the width and height of an image respectively. Then,

the patch W whose internal graphlets' joint probability can be maximized is selected, *i.e.*,

$$\max_W \prod_{G \in W} p(y(G) \rightarrow \mathbf{u}_d). \quad (12)$$

Patch W indicates the visual attribute that is most responsive to textual attribute $\mathbf{u}_d$. It captures the aesthetically relevant visual feature and we thus call it aesthlet.

C. Aesthlet-normalized CNN for Aesthetic Modeling

We integrate aesthlets into a deep architecture which learns patch-normalized representations to model visual attributes. After localizing different aesthlets in a weakly supervised way, we normalize and feed them into a CNN to extract standard representations for visual attributes modeling. To this end, we leverage both the power of CNN to learn discriminative visual features and the advantage of aesthlets to enhance the CNN learning by localizing highly aesthetic regions. Noticeably, although CNN has been successfully applied in a variety of multimedia tasks, they cannot generalize well when are trained using small-scale data. Our proposed aesthlets make the learning process require fewer training examples due to the increased training set size. This is because a Flickr image usually contains multiple aesthlets, each of which is treated as a single training example.

Starting from a large collection of aesthlet patches, we resize each of them into 64×64 pixels. Then, we randomly jetter each patch and flip it horizontally/vertically with probability 0.5 to improve the generalization, thereby a CNN is trained to represent each aesthlet. The architecture of the proposed CNN is elaborated in Fig. 4. The network contains four layers, *i.e.*, convolution, max pooling, local response normalization and a fully-connected layer with 1024 hidden units. After this, the network branches out one fully connected layer containing 128 units to describe each visual attribute. The last two layers are split to form tailored features for each visual attribute, *e.g.*, determining whether a Flickr image is "brightly-colored" or "well-structured" may require different visual features. Comparatively, the bottom layers are shared in order to: 1) reduce the number of parameters, and 2) take advantage of the common low-layer CNN structure. The parameters of our developed CNN are adjusted by cross-validation. Specifically, we first set the entire parameters: the input patch size, the number of convolutions, and strides exactly the same as those in Krizhevsky *et al.* [24]'s work. Then, we tune one parameter while leaving the rest unchanged. We set the tuning parameter by maximizing the accuracy of aesthetic prediction on our compiled data set [35].

The entire CNN is trained based on the standard back-propagation [16] of the error, combined with a stochastic gradient decent as a loss function, *i.e.*, the sum of the log-loss of each aesthlet from a training image. The architecture of each CNN layer is shown in Fig. 4 and the implementation details are elaborated in [16]. To effectively handle noise, we prone to employ aesthlet patches with high detection accuracies (based on (12)) in the training stage. In our paper, we did not use the pre-trained network. Actually, in our experiment, similar to [24], we have tried to adopt a pre-trained network (trained from three aesthetic images data set: CHUK [33], PNE [3], and AVA [2] respectively) to increase the aesthetic prediction accuracy.

As shown in Fig. 5, each aesthlet describes the aesthetics

Fig. 5. Examples comparing the roles of global composition and aesthlets in describing image aesthetics. The color-tagged regions indicate aesthlets calculated based on the Flickr tags on the right.

of a Flickr image from a single view. Therefore, we first utilize the developed CNN with respect to each aesthlet to generate representative aesthetic feature from a single view. Then, we combine the representations from multiple views to obtain the final aesthetic feature describing the entire image. More specifically, we extract the activations from the fc_attr layer in Fig. 4, which is 1024-dimensional, to describe each aesthlet. Finally, we concatenate the activations of all the aesthlets into a feature vector. It is worth emphasizing that, if an aesthlet does not activate for the image, we simply set the corresponding feature representation to zero.

As exemplified in Fig. 5, the aforementioned aesthlet-based CNN exploits regional visual features perceived by human beings. But the global features also plays an important role in describing image aesthetics. However, typically aesthlet patches cannot cover the entire image region. Even worse, in some degenerated cases, a Flickr image may have very few aesthlets detected, *e.g.*, abstract paintings without specific objects. To deal with this, we also incorporate a CNN whose inputs are patches covering the whole Flickr image, in order to capture the global aesthetics of Flickr images. Obviously, patches covering the whole image can be considered as special aesthlets. Therefore, the CNN is implemented with the same structure as that shown in Fig. 4.

Lastly, we concatenate the feature vectors corresponding to all the aesthlets from each Flickr image into a $128 * (D + 1)$-dimensional feature vector, which reflects the aesthetic feature of the image both locally and globally.

D. Applications of the Deep Aesthetic Features

We can employ the learned deep aesthetic features to enhance three applications in multimedia field: photo retargeting, aesthetics-based image classification and retrieval.

Image retargeting: For image retargeting, a 5-component GMM is adopted to learn the distribution of deep features calculated from all training well-aesthetic Flickr images:

$$p(z|\theta) = \sum_{l=1}^{5} \alpha_l \mathcal{N}(z|\pi_l, \Sigma_l), \quad (13)$$

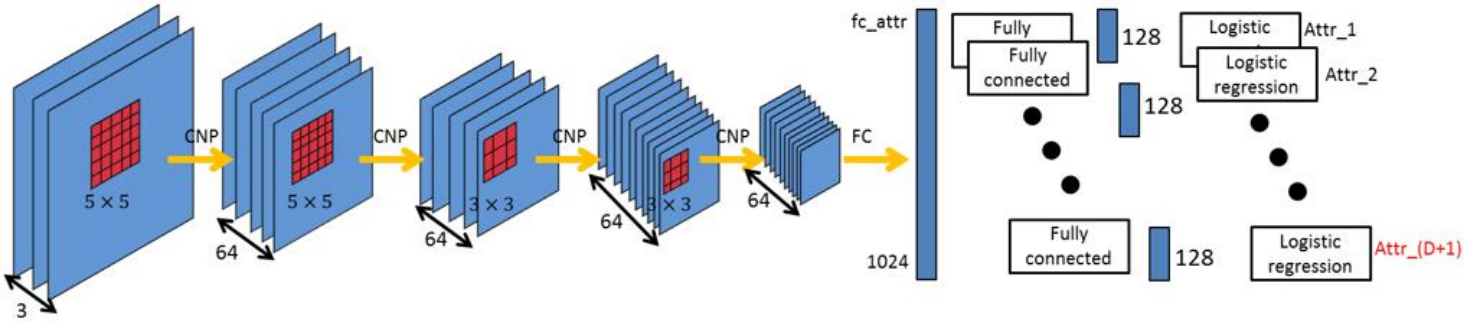

Fig. 4. A graphical illustration of the proposed deep architecture (CNP means convolution, normalization and pooling, and FC denotes the fully-connected layer. The $(D+1)$-th visual attribute corresponds to the aesthlet describing an entire Flickr image.)

where z denotes the deeply-learned features, and $\theta = \{\alpha_l, \pi_l, \Sigma_l\}$ are the GMM parameters.

Based on the GMM, we shrink a test image to make its deep feature most similar to those from the training images. Particularly, we decompose an image into equal-sized grids. Then, the horizontal (w.r.t. vertical) weight of grid ϕ is calculated as:

$$w_h(\phi) = \max p(z(\phi)|\theta), \tag{14}$$

where $z(\phi)$ denotes the deep aesthetic feature calculated based on the shrunk image. After obtaining the horizontal (w.r.t. vertical) weight of each grid, a normalization step is carried out to make them sum to one:

$$\bar{w}_h(\phi_i) = \frac{w_h\phi_i}{\sum_i w_h(\phi_i)}, \tag{15}$$

Given the size of the retargeted image $W \times H$, the horizontal dimension of the i-th grid is shrunk to $[W \cdot \bar{w}_h(\phi_i)]$, where $[\cdot]$ rounds a real number to the nearest integer. The shrinking operation along the vertical direction is similar to that along the horizontal direction.

Aesthetics-based image classification/retrieval: The deep aesthetic feature reflects human aesthetic perception. They can be used to: 1) identify whether a Flickr image is highly or low aesthetic, and 2) aesthetics-based image retrieval. The key of these two applications is a probabilistic model.

As shown in Fig. 6, given a set of training images and a test one, they are highly correlated through their deep aesthetic features z and z^* respectively. The probabilistic model contains four layers. The first layer corresponds to all the training images $I^1, I^1, \cdots, I^L$ which are aesthetically pleasing; the second layer denotes all the deep aesthetic features z learned from the training images; the third layer represents the deep aesthetic feature z_* extracted from the test image; and the last layer denotes the test image I_*.

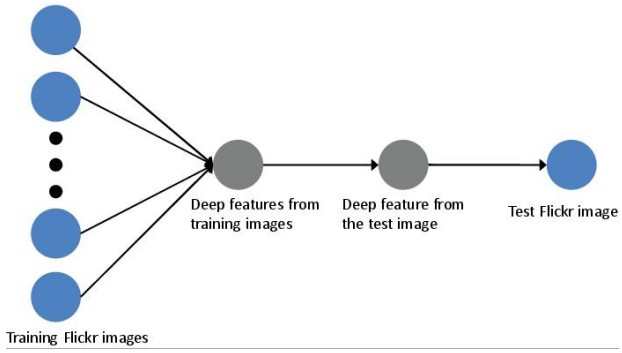

Fig. 6. A graphical illustration of the probabilistic model for image aesthetics quantification

Naturally, image aesthetics can be quantified as the amounts of deep aesthetic features that can be transferred from the training images into the test one. Thus, the aesthetics of a test image can be quantified as:

$$\begin{aligned}\gamma &= p(I_*|I^1, I^1, \cdots, I^L) \\ &= p(I_*|z_*) \cdot p(z_*|z) \cdot p(z|I^1, I^2, \cdots, I^L),\end{aligned} \tag{16}$$

where the probability $p(z_*|z)$ is calculated as: $p(z_*|z) = \prod_{j=1}^{L} p(z_*|z^j)$. z_* is the deep aesthetic feature calculated from the test image. z^j denote the deep aesthetic feature calculated from the j-th training image.

Following many algorithms [32], [35], we define the similarity between deep aesthetic features as a Gaussian kernel:

$$p(z_*|z) \propto \exp\left(-\frac{\|z_* - z\|^2}{2\sigma^2}\right). \tag{17}$$

After obtaining the aesthetic score γ of a test Flickr image. If $\gamma > 0.5$, then this image is deemed as "aesthetically pleasing", and vice versa. Besides, for aesthetics-based image retrieval, we output images in the database whose aesthetic scores are similar to that of the query image.

IV. Experimental Results and Analysis

This section evaluates the performance of the proposed deep aesthetic feature based on four experiments. The first experiment visualizes and analyzes the effectiveness of the proposed aesthlet. Then, we compare the three applications based on our deeply-learned aesthetic feature with the state-of-the-art. A step-by-step evaluation of the proposed method is presented subsequently. The last experiment evaluates the influence of different parameter settings.

All the experiments were carried out on a personal computer equipped with an Intel i5-2520M CPU and 8GB RAM. The algorithm was implemented on the Matlab 2012 platform.

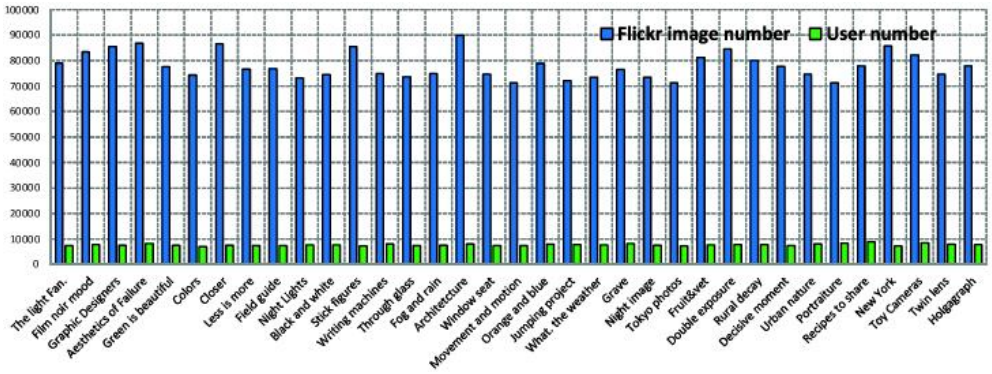

Fig. 7. The 35 Flickr groups (the horizontal axis) and the number of users (the vertical axis) in each of group

A. Descriptiveness of the Proposed Aesthlet

As the input to our deep architecture, aesthlets are image patches which are aesthetically pleasing and correspond to the textual attributes in each Flickr image. In this experiment, we first visualize the extracted aesthlets on a subset of our own complied Flickr image [34] dataset. Then, a comprehensive user study based on 132 participants is conducted to evaluate the descriptiveness of our proposed aesthlets.

Dataset Compilation: We spent significant time, effort and resources to crawl photos from 35 well-known groups form Flickr. For each group, we collected $70,000 \sim 90,000$ photos from nearly 7,400 Flickr users. The statistics of our dataset is shown in Fig. 7. For different Flickr groups, the numbers of photos belonging to each user varies from 10 to 220. We rank the Flickr images from each group based on the aesthetic measure by Zhang *et al.* [35]. The top 10% highly aesthetic photos constitute the image set for our experiment.

We set the number of textual attributes $D = 10$ and calculate the corresponding aesthlets according to (12). The representative aesthlets describing each textural attribute is presented in Fig. 8. The following observations can be made:

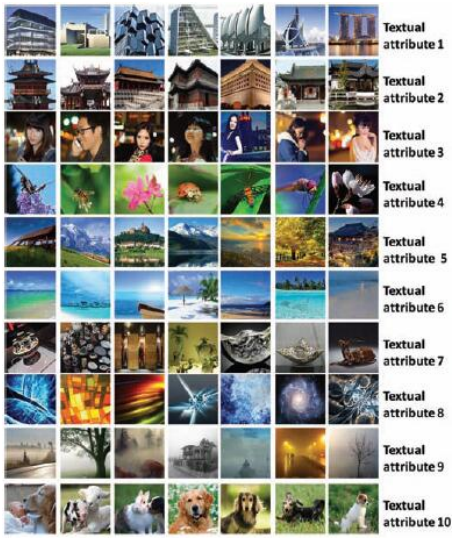

Fig. 8. Visualized aesthlets corresponding to each textual attribute

- The learned textural attributes are representative to each Flickr image, as the corresponding visual attributes can accurately detect aesthetically pleasing regions in each image. Moreover, we notice that each Flickr image is associated with fewer than three textural attributes. This is achieved by the sparse LSA and can accelerate the aesthetic modeling remarkably.
- Our graphlet-based weakly supervised learning algorithm effectively maps the textual attributes to the corresponding aesthelts. The is because graphelts can seamlessly capture objects with various shapes, and the weakly supervised algorithm can be solved analytically and efficiently.
- We notice that aesthlets corresponding to the same textual attribute have highly similar semantics, such as "modern architecture". Simultaneously, aesthlets corresponding to different textual attributes have distinguishable semantics, such as "scenery" and "beach". This demonstrates the effectiveness of our aesthlets discovering mechanism, since the ratio between inter-attribute scatter and intra-attribute scatter is prone to be maximized.

As aesthlets reflect human aesthetic perception, it is infeasible to measure their descriptiveness quantitatively. In this experiment, we conduct a user study to evaluate them qualitatively. This strategy was also adopted in [35] for comparing the aesthetic quality of multiple cropped photos. We invited 132 participants and most of them are master/Phd students from computer sciences department. For each aesthlet, we asked each participant to indicate its preference and discrimination scores. The former reflects the attractiveness of this aesthlet, while the latter shows whether this aesthelt captures a unique type of aesthetics. We set of number of textual attributes D to 10, 20, 40, 80 and 160 respectively. Then, the average preference and discrimination scores of aesthlets are calculated and reported in Table I. As can be seen, the preference scores under different values of D are above 0.9, which shows that the extracted aesthlets can accurately localizing visually attractive regions. Comparatively, the discrimination of aesthlet decreases dramatically when the value of D increases. The reason might be that the number of latent aesthetic types is much smaller than the number of textual attributes D.

B. Comparative Study

This subsection evaluates the three applications based on our deeply-learned aesthetic features: image retargeting,

Fig. 9. Photos retargeted based on different algorithms

TABLE I
AVERAGE PREFERENCE AND DISCRIMINATION SCORES OF AESTHLETS BY VARYING D

	$D = 10$	$D = 20$	$D = 40$	$D = 80$	$D = 160$
Preference	0.9434	0.9221	0.9121	0.9098	0.9002
Discrimination	0.9432	0.8956	0.8211	0.7676	0.6545

aesthetics-based image classification and retrieval.

Image retargeting: Fig. 9 compares the proposed method (PM) against several representative state-of-the-art approaches, including: seam carving (SC) [38] and its improved version (ISC) [37], optimized scale-and-sketch (OSS) [39] and saliency-based mesh parametrization (SMP) [40], and the patch-based wrapping (PW) [41]. Resolution of the resulting images is fixed to 640×960. The experimental images are all from the RetargetMe dataset [42]

In order to make the evaluation comprehensive, we adopt a paired comparison-based user study to evaluate the effectiveness of our proposed retargeting method. In the paired comparison, each participant is presented with a pair of retargeted photos from two different approaches, and required to indicate a preference as of which one they would choose for an iPhone wallpaper. In our user study, the participants are 40 amateur/professional photographers. As shown in Fig. 9, compared with its competitors, our approach well preserves the semantically important objects in the original photos, such as the warship and the bubble car. In contrast, the compared retargeting methods sometimes shrink the semantically important objects, such as the viewer, and boy and the football player. Even worse, SC and its variant ISC, as well as OSS may result in visual distortions.

In addition, we present an in-depth analysis of the user study on the 12 sets of retargeted photos in Fig. 9, which is inspired by the comparative study of retargeting algorithms in [42]. We invited 132 volunteers and most of them are master/Phd students from computer sciences department. For each volunteer, we asked them two questions. 1) Whether each visualized aesthlet is preferred by him/her and further indicate its aesthetic level, which is a real number ranging from 0 to 1. 0 indicates the lowest aesthetics while 1 denotes the highest aesthetics. 2) Whether each visualized aesthlet corresponding to a textual attribute is distinguishable from those corresponding to the rest textural attribute, and further indicate its discriminative level. 0 indicates the lowest discrimination while 1 denotes the highest discrimination. After collecting the preference and discrimination scores of all the aesthlets, we average them. Obviously, these two measures can show the effectiveness of our proposed method. First, we evaluate the degree of agreement when the volunteers vote for their favorite retargeted images, where a high disagreement reflects the difficulty in decision making. In our experiment, we use the coefficient of agreement defined by Kendall and Babington-Smith [43]. The coefficients over all the 12 retargeted images are shown in the last column of Fig. 10(a). Besides, we also collect the volunteers' votes on each attribute of the four sets of retargeted photos. As shown from the second column to the seventh column in Fig. 10(a), the volunteers are highly agreable on the face/people, the texture, and the symmetry because these attributes are well preserved by our retargeting model. Then, we present the votes on each attribute based on the 12 retargeted images. As shown in Fig. 10(b), the proposed method receives the most votes on all the attributes consistently.

Aesthetics-based image classification: We compare our approach with five image aesthetics evaluation methods. The compared methods include three global feature-based approaches proposed by Dhar *et al.* [5], Luo *et al.* [6] and Marchesotti *et al.* [7] respectively, two local patch integration-based methods proposed by Cheng *et al.* [36] and Nishiyama *et al.* [10] respectively, and the CNN-based aesthetic model by Lu *et al.*. In [22], Lu *et al.* proposed a novel framework to predict image style, aesthetics, and quality. The key technique is a deep network which learns the fine-grained details from

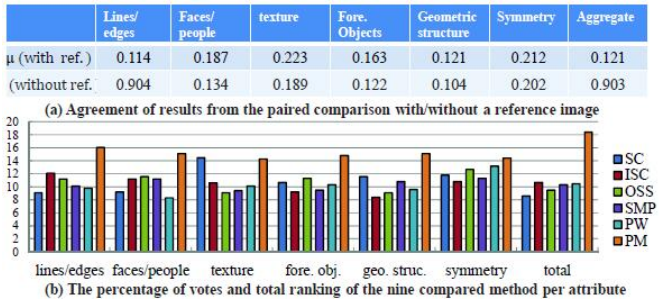

Fig. 10. A detailed analysis of the comparative retargeted photos in Fig. 9

TABLE II
COMPARATIVE AESTHETIC QUALITY PREDICTION ACCURACIES

	CUHK	PNE	AVA
Dhar *et al.*	0.7386	0.6754	0.6435
Luo *et al.*	0.8004	0.7213	0.6879
Marchesotti *et al.* (FV-Color-SP)	0.8767	0.8114	0.7891
Cheng *et al.*	0.8432	0.7754	0.8121
Nishiyama *et al.*	0.7745	0.7341	0.7659
Lu *et al.*	0.8043	0.8224	0.8213
The proposed method	**0.8879**	**0.8622**	**0.8465**

multiple image patches, where multi-patch aggregation functions can be learned as part of neural network training. Our proposed aesthlet-based image aesthetic model is significantly different from Lu *et al.*'s method, which is conducted in an unsupervised way. In our approach, the image patches (*i.e.*, aesthlets) are detected using a weakly-supervised learning algorithm, wherein the weak labels are latent topics detected from the M-dimensional augmented frequency vector.

In the comparative study, we observe that the source codes of the five baseline methods are unavailable and some experimental details are not provided. This makes it difficult to implement them exactly. In our implementation, we try to strengthen some components of the baseline methods. The following settings are employed. For Dhar's approach, we use the public codes from Li *et al.* [47] to extract the attributes from each photo. These attributes are combined with the low-level features proposed by Yeh *et al.* [48] to train the aesthetic classifier. For Luo *et al.*'s approach, not only are the low-level and high-level features in their publication implemented, but also the six global features from Getlter *et al.* [49]'s paper are utilized to strengthen the aesthetic prediction ability. For Marchesotti *et al.*'s approach, similar to the implementation of Luo *et al.*'s method, the six additional features are also adopted. For Cheng *et al.*'s algorithm, we implement it as a reduced version of our method, where only 2-sized graphlets are utilized for aesthetics measure. It is worth emphasizing that, for the probabilistic aesthetics evaluation models proposed by Cheng *et al.*, Nishiyama *et al.* method, and us, if the aesthetic score is higher than 0.5, then this image is categorized as highly aesthetic, and vice versa.

We present the aesthetics prediction accuracies on the CUHK, PNE and AVA in Table II. The image tags from the three datasets are assigned by ourselves. We hired 50 master/Phd students from our department to conduct the label annotation task. Each student spent 2 $\sim$ 4 hours to annotate 500 $\sim$ 700 images. All the images from the three data sets: the CHUK (12,000 images), the PNE (nearly 1700 images), and the AVA (about 25,000 images) are annotated. As shown, our approach outperforms Marchesotti *et al.*'s approach by nearly 2%, and exceeds the rest of the compared methods by over 6%, The results demonstrate the clear advantage of our method.

Aesthetics-guided image retrieval: We adopt the precision rate [52] to evaluate the performance of image retrieval based on our deeply-learned aesthetic features. Precision denotes the ratio of the number of relevant images (to the user) to the scope S, which is specified by the number of top-ranked images. In the experiment, we observe that it is sufficient to display 30 retrieved images on a screen. Presenting more images may decrease the quality of the presented images. Therefore, we set $S = 30$ throughout the experiment. The experimental dataset is our own crawled large-scale Flickr images from 35 groups.

In the current image retrieval systems, typically the query image is not exist in the image database. To handle this problem, we randomly select 30 images from each Flickr group as query images, while the rest are treated as the database for retrieval. Notably, the precision rate is calculated by averaging the results over the $30 * 35 = 1050$ queries.

To alleviate computational burden, we select the most informative images from the top 400 images. Relevance feedback is adopted to bridge the semantic gap in the retrieval system. The relevance feedback system contains a two-stage process: i) users are shown 30 images deemed most informative by SVM_{active}; ii) users give relevance feedback; iii) (i) is repeated taking into account feedback from ii). Users are asked to make relevance judgements toward the queries results. Thereafter, the feedback information is utilized to re-rank the images in the database. SVM_{active} is used as the relevance feedback algorithm. It provides users with the most informative images with respect to the ranking function.

We compare our deep features with well-known aesthetic descriptors proposed by Marchesotti *et al.* [7], Cheng *et al.* [36], Lu *et al.* [22] and Champbell *et al.* [23] respectively. As shown in Table III, in most of the 35 Flickr groups, our deeply-learned aesthetic feature achieves the best precision. We also observe that, although Marchesotti *et al.*'s aesthetic features is simple, its performance is competitive.

C. Step-by-step Model Justification

This experiment validates the effectiveness of the three key components in our deep aesthetic feature learning framework: 1) sparsity-constrained textual attributes discovery; 2) weakly supervised visual attributes localization; and 3) the aesthlet-normalized CNN training. To empirically demonstrate the effectiveness and inseparability of these components, we replace each component by a functionally reduced counterpart and report the corresponding performance. We focus on the application of aesthetics-based image classification.

Step 1: To demonstrate the usefulness of the sparse textual discovery, three experimental settings are used to weaken the adopted sparse LSA. First, we abandon the sparsity constraint in (4). Afterward, we replace the sparse LSA by the well-known linear discriminate analysis (LDA) and

TABLE III
PRECISION AT TOP 30 RETURNS OF THE FIVE COMPARED AESTHETIC FEATURES. THE HIGHEST PRECISION IS IN BOLD FOR EACH FLICKR GROUP.

Flickr group	March.	Cheng	Lu	Champ.	Ours
The light Fan.	0.34	0.33	0.43	0.31	**0.53**
Film noir Mood	0.16	0.19	0.23	0.16	**0.26**
Graphic designers	**0.33**	0.27	0.31	0.24	0.29
Aesthetics failure	0.11	0.09	0.15	0.12	**0.19**
Green is beautiful	0.14	0.13	0.16	0.12	**0.22**
Colors	**0.34**	0.25	0.28	0.21	0.31
Closer	0.43	0.41	0.45	0.36	**0.51**
Less is more	0.26	0.22	0.25	0.21	**0.34**
Field guide	0.34	0.31	0.33	0.28	**0.41**
Night lights	0.16	0.14	0.16	0.13	**0.19**
Black and white	**0.21**	0.18	0.19	0.15	0.18
Stick figure	0.43	0.41	0.34	0.37	**0.47**
Writing mach.	0.67	0.64	0.68	0.56	**0.76**
Through glass	0.09	0.09	0.07	0.05	**0.12**
Fog and rain	0.21	0.20	0.18	0.19	**0.25**
Architecture	0.54	0.53	0.54	0.46	**0.59**
Window seat	0.42	0.39	0.40	0.36	**0.46**
Movement	0.35	0.31	0.33	0.31	**0.37**
Orange and blue	0.18	0.16	0.18	0.13	**0.21**
Jump Project	0.26	0.24	0.19	0.22	**0.31**
What. the weather	0.14	0.12	0.11	0.10	**0.17**
Grave	0.28	0.24	0.26	0.21	**0.31**
Night images	0.14	0.16	0.12	0.11	**0.19**
Tokyo photos	0.24	0.21	0.19	0.19	**0.26**
Fruit&vet	0.65	0.62	0.61	0.62	**0.69**
Double exposure	0.28	0.23	0.24	0.19	**0.32**
Rural decay	0.25	0.23	0.24	0.19	**0.28**
Decisive moment	0.24	0.22	0.19	0.17	**0.27**
Urban nature	**0.32**	0.26	0.27	0.23	0.29
Portraiture	0.54	0.46	0.48	0.47	**0.59**
Recipes to share	0.26	0.22	0.24	0.21	**0.31**
New York	0.27	0.23	0.25	0.21	**0.31**
Toy cameras	0.37	0.31	0.31	0.30	**0.42**
Twin lens	0.26	0.22	0.24	0.21	**0.29**
Holgagraph	0.33	0.28	0.25	0.21	**0.38**
Average	0.28	0.23	0.21	0.19	**0.33**

principle component analysis (PCA) respectively. We present the calculated image classification accuracies in Table IV. As shown, removing the sparsity constraint results in an accuracy decrement of 5.2% on average. Moreover, either LDA nor PCA cannot optimally discover the latent textual attributes, since their performances lag behind ours by over 10%.

Step 2: To evaluate the effectiveness of the weakly supervised visual attributes learning, three experimental settings are adopted. We first replace the graphlet-based object detection by the objectness measure [53] and the part-based object detector [54] respectively. Then, we replace the superpixel-based spatial pyramid by the standard grid-based one. As shown in Table IV, on average, both the objectness and part-based detector cause a nearly 1.2% accuracy decrement. Noticeably, the training process of them requires manually annotated windows, which might be computationally intractable. Besides, grid-based spatial pyramid severely decrease the performance of aesthetics-based image classification. The reason is, compared with the superpixel-based graphlets, grid-based graphlets cannot fit the various object shapes optimally.

Step 3: We study the performance of our developed CNN. We remove the fully-connected layer and observe that the aesthetics-based classification reduces by 4% on average. This

TABLE IV
PERFORMANCE DECREMENT BY VARYING THE EXPERIMENTAL SETTINGS IN EACH COMPONENT

	CUHK	PNE	AVA
Step 1: Sparsity	5.4%	4.9%	5.2%
Step 1: Remove LSA	7.5%	6.7%	7.1%
Step 1: LDA	9.6%	10.9%	10.3%
Step 1: PCA	11.3%	9.5%	10.7%
Step 2: Objectness	1.3%	1.0%	1.1%
Step 2: Part-based	1.1%	1.1%	1.4%
Step 2: Superpixel	7.9%	8.4%	7.3%
Step 3: Fully-connected	4.3%	3.8%	4.0%
Step 3: Remove global descriptor	8.4%	6.2%	6.7%
Step 3: Only global descriptor	12.3%	14.1%	12.7%

result reflects the necessity to explore the common low-layer CNN structure. Additionally, we abandon the global image visual attribute. On average, a decrement of nearly 7% is observed. This shows the importance of explicitly modeling global image configurations in aesthetic modeling.

D. Parameter Analysis

In this experiment, we evaluate the influence of important parameters in our proposed aesthetic model: 1) the number of textual attributes D, 2) the regularizer weight λ in sparse LSA and 3) the structure of our designed CNN.

First, we tune the value of D from five to 80 with a step of

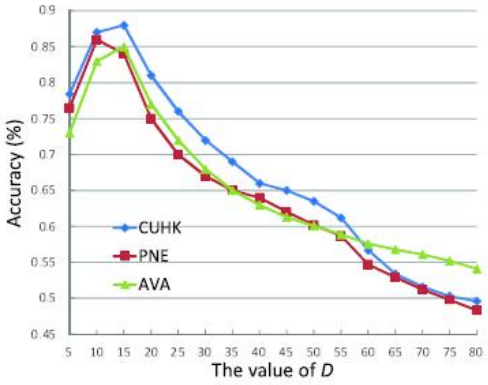

Fig. 11. Aesthetics-based image classification accuracies by varying D on the CUHK, PNE and AVA datasets

five and report the corresponding performance of aesthetics-based image classification. As shown in Fig. 11, on all the three datasets, the best accuracies are achieved when D is set to 10 or 15. This indicates there are about $10 \sim 15$ latent semantic categories from each of the three datasets. Second, we choose the value of λ from $\{0.5, 0.2, 0.1, 0.05, 0.001\}$ and report the corresponding accuracies. As shown in Table V, we notice that the best accuracies on all the three datasets are achieved when $\lambda = 0.1$. Next, we testify the effectiveness of our developed five-layered CNN. We change our CNN to a four- and six-layered CNN respectively. It is noticeable that the aesthetics-based image classification accuracies decreased by 4.3% and 6.7% respectively. Actually, in our implementation, the five CNN layers are validated by cross validations. Lastly, we preserve only the global descriptor for image aesthetics

TABLE V
AESTHETICS-BASED IMAGE CLASSIFICATION ACCURACIES UNDER DIFFERENT VALUES OF λ

Dataset	$\lambda = 0.5$	$\lambda = 0.2$	$\lambda = 0.1$	$\lambda = 0.05$	$\lambda = 0.001$
CHUK	0.7612	0.8133	0.8879	0.6454	0.7687
PNE	0.7453	0.8231	0.8622	0.7121	0.7453
AVA	0.8113	0.8214	0.8465	0.7376	0.7786

prediction, where the CNN-based descriptor is abandoned. As shown in the last row of Table V, aesthetics-based image classification using only global descriptor results in an accuracy decrement of 13% on average. This observation clearly demonstrates the necessity of exploiting local descriptors in aesthetic modeling.

V. CONCLUSIONS AND FUTURE WORK

Perceptually aesthetic model is an important topic in multimedia field [44], [45], [46]. This paper proposes a CNN framework to hierarchically model how humans perceive aesthetically pleasing regions in a Flickr image. We first calculate a compact set of textual attributes from tagged Flickr images using a sparsity-constrained LSA. Then, a weakly supervised learning paradigm projects the textual attributes onto the corresponding aesthlets in each image. These aesthlets capture visually attractive image regions and are deployed to train a CNN to mimick human aesthetic perception. Based on the CNN, we represent each Flickr image by a set of deeply-learned aesthetic features, which can enhance a series of multimedia applications, *e.g.*, image retargeting, aesthetics-based image classification and retrieval.

In the future, this work will be extended to a more comprehensive deep architecture that encodes auxiliary visual cues such as exposure, contrast and symmetry, which might also be contributive to image aesthetic modeling.

REFERENCES

[1] Thomas Hofmann, Probabilistic Latent Semantic Analysis, *in Proc. of UAI*, 1999.

[2] Luca Marchesotti, Naila Murray, Florent Perronnin, Discovering Beautiful Attributes for Aesthetic Image Analysis, *IJCV*, 10.1007/s11263-014-0789-2, 2014.

[3] Ritendra Datta, Dhiraj Joshi, Jia Li, James Z. Wang, Studying Aesthetics in Photographic Images using a Computational Approach, *in Proc. of ECCV*, 2006.

[4] Yan Ke, Xiaoou Tang, Feng Jing, The Design of High-level Features for Photo Quality Assessment, *in Proc. of CVPR*, 2006.

[5] Sagnik Dhar, Vicente Ordonez, Tamara L. Berg, High Level Describable Attributes for Predicting Aesthetics and Interestingness, *in Proc. of CVPR*, 2011.

[6] Wei Luo, Xiaogang Wang, Xiaoou Tang, Content-based Photo Quality Assessment, *in Proc. of ICCV*, 2011.

[7] Luca Marchesotti, Florent Perronnin, Diane Larlus, Gabriela Csurka, Assessing the Aesthetic Quality of Photographs using Generic Image Descriptors, *in Proc. of ICCV*, 2011.

[8] Bin Cheng, Bingbing Ni, Shuicheng Yan, Qi Tian, Learning to Photograph, *ACM MM*, 2010.

[9] Masashi Nishiyama, Takahiro Okabe, Yoichi Sato, Imari Sato, Sensation-based Photo Cropping, *ACM MM*, 2009.

[10] Masashi Nishiyama, Takahiro Okabe1, Imari Sato, Yoichi Sato, Aesthetic Quality Classification of Photographs based on Color Harmony, *in Proc. of CVPR*, 2011.

[11] Subhabrata Bhattacharya, Rahul Sukthankar, Mubarak Shah, A Framework for Photo-Quality Assessment and Enhancement based on Visual Aesthetics, *ACM MM*, 2010.

[12] Luming Zhang, Yue Gao, Roger Zimmermann, Qi Tian, Xuelong Li, Fusion of Multi-Channel Local and Global Structural Cues for Photo Aesthetics Evaluation, *IEEE T-IP*, 23(3): 1419–1429, 2014.

[13] Luming Zhang, Yue Gao, Rongrong Ji, Qionghai Dai, Xuelong Li, Actively Learning Human Gaze Shifting Paths for Photo Cropping, *IEEE T-IP*, 23(5): 2235C2245, 2014.

[14] Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, Stefan Carlsson, CNN Features off-the-shelf: an Astounding Baseline for Recognition, *arXiv preprint arXiv:1403.6382*, 2014.

[15] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, Trevor Darrell, DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition, *in Proc. of ICML*, 2014.

[16] Alex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton, ImageNet Classification with Deep Convolutional Neural Networks, *in Proc. of NIPS*, 2012.

[17] Ning Zhang, Manohar Paluri, MarcAurelio Ranzato, Trevor Darrell, Lubomir Bourdev, PANDA: Pose Aligned Networks for Deep Attribute Modeling, *in Proc. of CVPR*, 2014.

[18] Ping Luo, Xiaogang Wang, Xiaoou Tang, A Deep Sum-Product Architecture for Robust Facial Attributes Analysis, *in Proc. of ICCV*, 2013.

[19] Ning Zhang, Jeff Donahue, Ross Girshick, Trevor Darrell, Part-based R-CNNs for Fine-grained Category Detection, *in Proc. of ECCV*, 2014.

[20] Pierre Sermanet, Koray Kavukcuoglu, Soumith Chintala, Pedestrian Detection with Unsupervised Multi-Stage Feature Learning, *in Proc. of CVPR*, 2014.

[21] Lin Sun, Kui Jiay, Tsung-Han Chany, Yuqiang Fang, Gang Wang, Shuicheng Yan, DL-SFA: Deeply-Learned Slow Feature Analysis for Action Recognition, *in Proc. of CVPR*, 2014.

[22] Xin Lu, Zhe Lin, Hailin Jin, Jianchao Yang, James Z. Wang, RAPID: Rating Pictorial Aesthetics using Deep Learning, *ACM MM*, 2014.

[23] Allan Campbell, Vic Ciesielksi, A. K. Qin, Feature Discovery by Deep Learning for Aesthetic Analysis of Evolved Abstract Images, *Evolutionary and Biologically Inspired Music, Sound, Art and Design Lecture Notes in Computer Science*, 9027: 27–38, 2015.

[24] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, Li Fei-Fei, ImageNet: A Large-Scale Hierarchical Image Database, *in Proc. of CVPR*, 2009.

[25] Scott Deerwester, Susan T. Dumais, George W. Furnas, Thomas K. Landauer, Richard Harshman, Indexing by Latent Semantic Analysis, *Journal of American Society for Information Sciences*, 41(6): 391–407, 1990.

[26] Xi Chen, Yanjun Qi, Bing Hai, Qihang Lin, Jaime G. Garbonell, Sparse Latent Semantic Analysis, *in Proc. of NIPS*, 2010.

[27] Robert Tibshirani, Regression Shrinkage and Selection via the Lasso, *Journal of the Royal Statistical Society: Series B*, 58: 267–288, 1996.

[28] Jerome Friedman, Trevor Hastie, Rob Tibshirani, Regularization Paths for Generalized Linear Models via Coordinate Descent, *Journal of Statistical Software*, 33(1): 2010.

[29] Markus Stricker, Markus Orengo, Similarity of Color Images, *Storage and Retrieval of Image and Video Databases*, 1995.

[30] Navneet Dalal, Bill Triggs, Histograms of Oriented Gradients for Human Detection, *in Proc. of CVPR*, 2005.

[31] Shiming Xiang, Feiping Nie, Yangqiu Song, Changshui Zhang, Chunxia Zhang, Embedding New Data Points for Manifold Learning via Coordinate Propagation, *Knowlege and Information Systems*, 19(2): 159C184, 2008.

[32] Mingli Song, Dacheng Tao, Chun Chen, Jiajun Bu, Jiebo Luo, Chengqi Zhang, Probabilistic Exposure Fusion, *IEEE T-IP*, 21(1): 341–357, 2012.

[33] Yan Ke, Xiaoou Tang, Feng Jing, The Design of High-level Features for Photo Quality Assessment, *in Proc. of CVPR*, 2006.

[34] Luming Zhang, Roger Zimmermann, Flickr Circles: Mining Socially-Aware Aesthetic Tendency, *in Proc. of ICME*, 2015.

[35] Luming Zhang, Mingli Song, Qi Zhao, Xiao Liu, Jiajun Bu, Chun Chen, Probabilistic Graphlet Transfer for Photo Cropping, *IEEE T-IP*, 21(5): 2887–2897, 2013.

[36] Bin Cheng, Bingbing Ni, Shuicheng Yan, Qi Tian, Learning to Photograph, *ACM MM*, 2010.

[37] Michael Rubinstein, Ariel Shamir, Shai Avidan, Improved Seam Carving for Video Retargeting, *ACM TOG*, 27(3), 16, 2008.

[38] Shai Avidan, Ariel Shamir, Seam Carving for Content-Aware Image Resizing, *ACM TOG*, 26(3),10, 2007.

[39] Yu-Shuen Wang, Chiew-Lan Tai, Olga Sorkine, Tong-Yee Lee, Optimized Scale-and-Stretch for Image Resizing, *ACM TOG*, 27(5), 118, 2008.

[40] Yanwen Guo, Feng Liu, Jian Shi, Zhihua Zhou, Michael Gleicher, Image Retargeting Using Mesh Parameterization, *IEEE T-MM*, 11(5): 856–867, 2009.
[41] Shih-Syun Lin, I-Cheng Yeh, Chao-Hung Lin, Tong-Yee Lee, Patch-Based Image Warping for Content-Aware Retargeting, *IEEE T-MM*, 15(2): 359–368, 2013.
[42] Michael Rubinstein, Diego Gutierrez, Olga Sorkine, Ariel Shamir, A Comparative Study of Image Retargeting, *ACM TOG*, 29(5),160, 2010.
[43] Michael Rubinstein, Ariel Shamir, Shai Avidan, Multi-operator Media Retargeting, *ACM TOG*, 28(3), 23, 2009.
[44] Liqiang Nie, Meng Wang, Yue Gao, Zheng-Jun Zha, Tat-Seng Chua, Beyond Text QA: Multimedia Answer Generation by Harvesting Web Information, *IEEE T-MM*, 15(2): 426-441, 2013.
[45] Zhigang Ma, Yi Yang, Nicu Sebe, Kai Zheng and Alexander G. Hauptmann, Multimedia Event Detection Using A Classifier-specific Intermediate Representation, *IEEE T-MM*, 15(7):1628–1637, 2013.
[46] Yi Yang, Jingkuan Song, Zi Huang, Zhigang Ma, Nicu Sebe and Alexander G. Hauptmann, Multi-Feature Fusion via Hierarchical Regression for Multimedia Analysis, *IEEE T-MM*, 15(3): 572–581, 2013.
[47] Fei-Fei Li, Pietro Perona, A Bayesian Hierarchical Model for Learning Natural Scene Categories, *in Proc. of CVPR*, 2005.
[48] Che-Hua Yeh, Yuan-Chen Ho, Brian A. Barsky, Ming Ouhyoung, Personalized Photograph Ranking and Selection System, *ACM MM*, 2010.
[49] Peter Gehler, Sebastian Nowozin, On Feature Combination for Multiclass Object Classification, *in Proc. of ICCV*, 2009.
[50] Luming Zhang, Meng Wang, Liqiang Nie, Liang Hong, Yong Rui, Qi Tian, Retargeting Semantically-Rich Photos, *IEEE T-MM*, 17(9): 1538–1549, 2009.
[51] Luming Zhang, Yue Gao, Richang Hong, Yuxing Hu, Rongrong Ji, Qionghai Dai, Probabilistic Skimlet Fusion for Summarizing Multiple Consumer Landmark Videos, *IEEE T-MM*, 71(1): 40-49, 2015.
[52] Dionysius P. Huijsmans, Nicu Sebe, How to Complete Performance Graphs in Content-based Image Retrieval: Add Generality and Normalize Scope, *IEEE T-PAMI*, 27(2): 245–251, 2005.
[53] Bogdan Alexe, Thomas Deselaers, Vittorio Ferrari, Measuring the Objectness of Image Windows, *IEEE T-PAMI*, 34(11): 2189-2202, 2012.
[54] Pedro F. Felzenszwalb, Ross B. Girshick, David A. McAllester, Deva Ramanan, Object Detection with Discriminatively Trained Part-Based Models, *IEEE T-PAMI*, 32(9): 1627-1645, 2010.